

Amdt. dated October 12, 2005
Reply to Office action of July 12, 2005

Serial No. 10/624,408
Docket No. P16578
Firm No. 0077.0025

REMARKS/ARGUMENTS

1. Amended Claims 18-27 Comply With 35 U.S.C. §101

The Examiner rejected claims 18-27 under 35 U.S.C. §101 (Section 101) as being directed to non-statutory subject matter for possibly covering the code in transmission lines and transmission signals. (Office Action, pg. 2) Applicants traverse this finding because encoding code or logic in a transmission signal, transmission line, etc. is just as tangible as encoding code in a magnetic or electronic storage media. The magnetic transitions and electrons that encode and represent code in a magnetic or electronic storage media are no more or less "tangible" than the signals and waves that encode and represent code being transmitted.

Nonetheless, to expedite prosecution, Applicants have amended claim 18 to recite that the "article of manufacture" comprises a device implementing code that causes the claimed operations to be performed to convince the Examiner to withdraw the Section 101 rejection and expedite prosecution. The Application on pg. 8, para. 18 discloses that the article of manufacture may comprise devices implementing code, such as computer readable medium, hardware logic, etc. Applicants make this amendment while continuing to traverse the legal basis for this Section 101 rejection, and reserve the right to pursue claims covering media that the Examiner has found to be intangible.

Applicants further added dependent claim 28 to recite that the device implementing the code may comprise a computer readable medium or a hardware component. The Application on page 8, para. 18 discloses that article of manufacture may comprise computer readable medium, hardware in which the code is embodied.

2. Claims 1, 10, 11, 17, 18, and 27 are Patentable Over the Cited Art

The Examiner rejected claims 1, 10, 11, 17, 18, and 27 as anticipated (35 U.S.C. §102(b)) by Foster (U.S. Patent No. 5,948,081). Applicants traverse for the following reasons.

Claims 1, 11, and 18 concern managing requests to an Input/Output (I/O) device, and require: queuing I/O requests directed to the I/O device; determining whether a number of queued I/O requests exceeds a threshold; if the number of queued I/O requests exceeds the threshold, then calculating a coalesce limit; coalescing a number of queued I/O requests not exceeding the calculated coalesce limit into a coalesced I/O request; and transmitting the coalesced I/O request.

Amdt. dated October 12, 2005
Reply to Office action of July 12, 2005

Serial No. 10/624,408
Docket No. P16578
Firm No. 0077.0025

The Examiner cited col. 7, lines 25-41 and the bridging paragraph between cols. 9 and 10 as disclosing these claim requirements. (Office Action, pg. 3)

The cited col. 7 mentions that read requests are removed from a queue before write requests and given higher priority than write requests as long as the number of write requests are below a threshold. When the pending write requests reaches the threshold, the queue logic asserts full to service write requests in lieu of read requests. From that moment on, requests are serviced from the write queue 46 and read queue 44 in a ping-pong fashion.

Nowhere does the cited col. 7 anywhere disclose the claim requirement of calculating a coalesce number if the number of queued I/O requests exceeds a threshold and then coalescing a number of queued I/O requests into a coalesced I/O request not exceed the calculated coalesce limit. There is no disclosure of calculating a coalesce limit of a number of I/O requests to coalesce from the queue and transmit. Instead, the cited col. 7 mentions that if a threshold is reached, then I/O requests are pulled from both the read and write queues. The cited col. 7 provides no mention or disclosure of disclosing the requirement of calculating a coalesce number indicating a number of I/O requests to be coalesced into a coalesce I/O request. Instead, in the cited col. 7, when the threshold is reached, write requests are then serviced.

The Examiner further referenced the "almost full condition" in col. 7. (Office Action, pg. 4) The cited "almost full condition" is asserted when the number of write requests exceed a threshold so that writes may be processed. However, nowhere is there any disclosure that the "almost full condition" is a coalesce limit that is calculated when the number of queued I/O requests exceeds a threshold indicating a number of requests to coalesce in a coalesced I/O request.

The cited cols. 9-10 mention comparing the number of memory requests removed from a requests queue with the number which are to be removed, the "N" number. However, it is not always possible to service the "N" number of requests in burst mode because some requests cannot be removed from the queue, such as if the requests do not have a zero count. If all three requests have a zero count, then all three can be removed. If the flush indicator is not zero, then the other queue is serviced first.

Nowhere do the cited cols. 9-10 anywhere disclose the claim requirement of calculating a coalesce number if the number of queued I/O requests exceeds a threshold and then coalescing the coalesced number of I/O requests into a coalesced I/O request. For instance, nowhere do the

Amdt. dated October 12, 2005
Reply to Office action of July 12, 2005

Serial No. 10/624,408
Docket No. P16578
Firm No. 0077.0025

cited cols. 9-10 anywhere disclose calculating a number of requests that will be transmitted, i.e., the coalesce number, if a threshold is reached. Instead, the cited cols. 9-10 discuss which queue to service bursts requests from, such that if a queue does not have a sufficient number of requests to service (Foster mentions 3 requests with a zero count), then the other queue is serviced. In other words, the cited cols. 9-10 nowhere disclose calculating the number of requests that will be coalesced in a burst request, where the coalesce number is calculated when the number of queued I/O requests exceed a threshold. For instance, the cited cols. 9-10 does not disclose calculating "N" if the number of queued I/O requests exceeds a threshold.

Accordingly, claims 1, 11, and 18 are patentable over the cited art because the cited Foster does not disclose all the claim requirements.

Claims 10, 17, and 27 are patentable over the cited art because they depend from one of claims 1, 11, and 18, which are patentable over the cited art for the reasons discussed above. Moreover, these claims provide additional grounds of patentability over the cited art for the following reasons.

Claims 10, 17, and 27 additionally recite transmitting one I/O request from the queue if the number of queued I/O requests does not exceed the threshold. The Examiner cited col. 9, lines 25-35 as disclosing the additional requirements of these claims. (Office Action, pg. 4) Applicants traverse.

The cited col. 9 mentions that if a read hit occurs, a read request having a similar address to the current write request must be serviced. However, if a read hit does not occur, the write can be serviced. After N bursts, then the write request queue is monitored to see if it is full. If the N number of burst has not been reached, then the bursts can continue for write or read requests.

Nowhere does the cited col. 9 anywhere disclose that one request is transmitted from a queue if the number of queued I/O requests in that queue does not exceed the threshold. Instead, the cited col. 9 discusses how a read request is serviced if there is a read hit.

Accordingly, claims 10, 17, and 27 provide additional grounds of patentability over the cited art because the cited art does not disclose the additional requirements of these claims.

3. Claims 2, 3, 5-9, 12, 13, 15, 16, 19, 20, and 22-26 are Patentable Over the Cited Art

The Examiner rejected claims 2, 3, 5-9, 12, 13, 15, 16, 19, 20, and 22-26 as obvious (35 U.S.C. §103(a)) over Foster in view of Myers (U.S. Patent No. 6,877,049). Applicants traverse.

Amdt. dated October 12, 2005
Reply to Office action of July 12, 2005

Serial No. 10/624,408
Docket No. P16578
Firm No. 0077.0025

First off, these claims are patentable over the cited art because they depend from one of claims 1, 11, and 18, which are patentable over the cited art for the reasons discussed above. Moreover the following dependent claims provide additional grounds of patentability over the cited art for the reasons discussed below.

Claims 2, 12, and 19 depend from claims 1, 11, and 18 and further require that the calculated coalesce limit dynamically varies based in part on the number of queued I/O requests. The Examiner cited col. 5, lines 48-65 of Myers as teaching the additional requirements of these claims. (Office Action, pg. 5) Applicants traverse.

The cited col. 5 discusses a credit counter that is incremented by a number of data blocks received from the memory to ensure that the credit value accounts for pending data requests that are still in process in the memory.

Although the cited Myers dynamically varies a credit counter indicating a number of pending requests still in process, nowhere does the cited Myers anywhere teach or suggest dynamically varying a coalesce limit that indicates a number of I/O requests to coalesce to transmit. Thus, Myers discusses dynamically varying a parameter, but a parameter that is different from the claimed coalesce limit.

Accordingly, claims 2, 12, and 19 provide additional grounds of patentability over the cited art because the cited art does not disclose the additional requirements of these claims.

Claims 3, 13, and 20 depend from claims 2, 12, and 19 and further require that calculating the coalesce limit includes dividing the number of queued I/O requests by an interval. The Examiner cited col. 6, lines 12-20 of Myers as teaching the additional requirements of these claims. (Office Action, pg. 5) Applicants traverse.

The cited col. 6 mentions a trigger threshold that is selected based on the rate by which data in the buffer is depleted by the client as well as the frequency by which data requests can be made. Nowhere does the cited col. 6 teach or suggest calculating a coalesce limit by dividing the number of queued I/O requests by an interval. Instead, the cited col. 6 mentions that the trigger threshold is based on a depletion rate as well as the frequency by which data requests can be made. There is no suggestion in the cited col. 6 of dividing the number of queued I/O requests by an interval as claimed.

Accordingly, claims 3, 13, and 20 provide additional grounds of patentability over the cited art because the cited art does not disclose the additional requirements of these claims.

Amdt. dated October 12, 2005
Reply to Office action of July 12, 2005

Serial No. 10/624,408
Docket No. P16578
Firm No. 0077.0025

Claims 5, 15, and 22 depend from claims 1, 11, and 18 and further require that I/O requests are queued in a first queue or a second queue, wherein determining whether the number of queued I/O requests exceeds the threshold comprises determining whether a number of I/O requests in the second queue exceeds the threshold, and wherein coalescing the number of queued I/O requests comprises coalescing I/O requests from the first queue.

The Examiner cited cols. 9-10 of Foster as teaching the claim requirement of determining whether a number of I/O requests in the second queue exceeds the threshold, and wherein coalescing the number of queued I/O requests comprises coalescing I/O requests from the first queue. (Office Action, pgs. 6-7)

The cited cols. 9-10 mention comparing the number of memory requests removed from a requests queue with the number which are to be removed, the "N" number. If all three requests have a zero count, then all three can be removed. If the flush indicator is not zero, then the other queue is serviced first.

Nowhere do the cited cols. 9-10 disclose that coalescing I/O requests in one queue if the number of I/O requests in a second queue exceeds a threshold. The cited cols. 9-10 mention that if all requests cannot be removed, then the other queue is serviced. However, this is servicing a different queue if one queue does not have enough N requests to burst. Nowhere do the cited cols. 9-10 teach or suggest that requests from one queue are coalesced and sent if the requests in another queue exceed a threshold. In fact, the cited cols. 9-10 appear to teach the opposite, because if not enough requests can be removed from the current queue, then the other queue is serviced. The claims coalesce requests to transmit from the first queue if the number of requests in the second queue exceed the threshold, not fall short.

Accordingly, claims 5, 15, and 22 provide additional grounds of patentability over the cited art because the cited art does not disclose the additional requirements of these claims.

Claims 6 and 23 depend from claims 5 and 22 and further require adding the transmitted coalesced I/O request to the second queue. The Examiner cited the above discussed cols. 9-10 of Foster as teaching the additional requirements of these claims. (Office Action, pgs. 6-7)
Applicants traverse.

The cited cols. 9-10 discuss servicing requests from a queue. However, when the requests from a queue are serviced, they are removed from the queue and transmitted in burst mode. Nowhere do the cited cols. 9-10 anywhere teach or suggest that requests from one of the queues

Amdt. dated October 12, 2005
Reply to Office action of July 12, 2005

Serial No. 10/624,408
Docket No. P16578
Firm No. 0077.0025

are added to the other queue. For the cited cols. 9-10, this would mean adding the requests from the read queue to the write queue or vice versa.

Further, nowhere does the cited art teach or suggest adding a coalesced number of requests from a first queue to a second queue, where the requests are added if the number of requests in the second queue exceeds a threshold.

Accordingly, claims 6 and 23 provide additional grounds of patentability over the cited art because the cited art does not disclose the additional requirements of these claims.

Claims 9, 16, and 26 depend from claims 5, 15, and 22 and further require determining whether there are at least two I/O requests in the first queue after determining that the number of requests in the second queue exceeds the first queue, wherein I/O requests from the first queue are only coalesced if there are at least two I/O requests in the first queue.

The Examiner cited Myers as teaching the claim requirement of determining whether there are at least two I/O requests in the first queue after determining that the number of requests in the second queue exceeds the first queue. (Office Action, pg. 7) The Examiner did not cite any particular section of Myers in making this finding. The Examiner previously cited col. 1, lines 40-51 and the bridging paragraph on cols. 4-5 of Myers with respect to two queues. (Office Action, pg. 6) Applicants submit that this cited Myers does not teach the claim requirements.

The cited col. 1 of Myers mentions a buffer that receives data requested from the memory that enables sufficient data to be immediately accessible to the client. Nowhere does this cited col. 1 anywhere teach or suggest determining the number of requests in a first queue after determining that the number of requests in the second queue exceeds the first queue. Instead, the cited col. 2 just discusses a buffer for memory.

The cited cols. 4-5 discusses a FIFO buffer, a credit value and trigger value used to determine when to make data requests upon the memory. Each credit represents a data request that can be made if the amount of data in the buffer falls to or below the trigger threshold. The credit value is incremented by the number of blocks received by the FIFO controller. The trigger value is yes if the threshold is exceeded. When the trigger value is yes, the controller sends a data request to the memory for a certain number of blocks. The trigger value is decremented by the number of data blocks requested.

Nowhere do the cited cols. 4-5 of Myers teach or suggest the claim requirement of determining whether there are at least two I/O requests in the first queue after determining that

Amdt. dated October 12, 2005
Reply to Office action of July 12, 2005

Serial No. 10/624,408
Docket No. P16578
Firm No. 0077.0025

the number of requests in the second queue exceeds the first queue. Instead, if the number of blocks received exceeds a threshold, a data request for more blocks is sent to the memory. The cited cols. 4-5 do not teach determining whether there are two requests in the first queue if the number of requests in the second queue exceeds the first queue. This specific claimed determination is not taught or suggested in the cited Myers.

Accordingly, claims 9, 16, and 26 provide additional grounds of patentability over the cited art because the cited art does not disclose the additional requirements of these claims.

4. Claims 4, 14, and 21 are Patentable Over the Cited Art

The Examiner rejected claims 4, 14, and 21 as obvious (35 U.S.C. §103(a)) over Foster in view of Gunlock (U.S. Patent No. 5,522,054). Applicants traverse.

Applicants submit that these claims are patentable over the cited art because they depend from claims 4, 14, and 21.

Conclusion

For all the above reasons, Applicant submits that the pending claims 1-27 are patentable over the art of record. Applicants have not added any claims. Nonetheless, should any additional fees be required, please charge Deposit Account No. 50-0585.

\\

The attorney of record invites the Examiner to contact him at (310) 553-7977 if the Examiner believes such contact would advance the prosecution of the case.

Dated: October 12, 2005

By: 

David W. Victor
Registration No. 39,867

Please direct all correspondences to:

David Victor
Konrad Raynes & Victor, LLP
315 South Beverly Drive, Ste. 210
Beverly Hills, CA 90212
Tel: 310-553-7977
Fax: 310-556-7984